

Daryl's TCP/IP Primer

Addressing and Subnetting on the Near Side of the 'Net

Contents

1. Disclaimer
2. Overview and Scope
3. The Bottom of the OSI Model
4. Intro to Ethernet
5. Why is IP so much more difficult than IPX?
6. IP Addresses, Subnet Masks, and Subnetting
7. Subnetting, Bit by Bit
8. Routing and Static Routes
9. Troubleshooting
10. TCP and UDP Communication
11. Network Address Translation (NAT)
12. The Domain Name System (DNS)
13. Tips for Building an IP LAN
14. WAN Connectivity
15. Update Notifications/Comment Form
16. Questions and Answers
17. Other Sources
18. Glossary

1. Disclaimer

This document is presented with no warranties or guarantees of ANY KIND including correctness or fitness for any particular purpose. The author(s) of this document have attempted to verify correctness of the data contained herein; however, slip-ups can and do happen. If you use this data, you do so at your own risk. This document is Copyright © 1996-1999 by Daryl Banttari, and is made available as a service to the Internet community. It may not be sold in any medium, including electronic, CD-ROM, or database, packaged with any commercial product, or published in print, without the explicit, written permission of Daryl Banttari. You may freely link or refer to this document at <http://ipprimer.windsorcs.com/> ; however, the author can make no guarantees of its future availability at that location. [I don't expect it to change, but I can't *guarantee* it won't change.]

If you register for update notification, I'll do my best to let you know of any changes in location in the future.

Finally, if you're in the Minneapolis (USA) area and want implementation help or training, drop me a note. My Resume.

2. Overview and Scope

This document is designed to give the reader a reasonable working knowledge of TCP/IP subnetting, addressing, and routing. It is not intended to be complete, or to cover all issues; I'm just tired of re-explaining this stuff, so now I can just point to this document instead of constantly generating 3-page emails :-). This is targeted toward LAN administrators just moving to TCP/IP, however it should help anyone who wants to know a little (more) about how TCP/IP works. This document does not, generally, apply to dial-up SLIP/PPP connections.

The difference between this (a primer) and an FAQ, is that most FAQ's, in practice, tend to be question-and-answer oriented, and generally seem to try to cover ALL issues, not just the ones frequently asked about. This primer is intended as a starting point for someone who has an interest in the subject, but doesn't know where to start or what questions to ask. This should also help to broaden the understanding of people who have worked with TCP/IP for a while, but either haven't had the time to study all the less-than-useful theory behind the subject, or have been somewhat overwhelmed by the many theoretical details and have missed the big picture.

This is HTML, but I have made it one large page for the benefit of those who prefer to print off a copy and read it that way. Also useful for sharing via hard copy. If you choose to print and distribute this, I ask that you distribute it in its entirety, and that you don't charge for it.

Feedback, of course, is always greatly appreciated, and will help determine the direction and growth of this living document. In fact, just a quick email to say "thanks" (if it helped) will help motivate me to keep this current and expanding :-)

3. The Bottom of the OSI Model

The OSI Networking Model is used as a reference point to describe how the various "layers" of networking interoperate. For this discussion, I will describe the bottom three layers:

Layer	Name	Protocols / Terms	Devices that operate in this layer	Addresses are called...
3	Network	IP, IPX, AppleTalk	Routers	Network Addresses
2	Datalink	Ethernet, Token Ring, PPP, SLIP, HDLC	Bridges, Switches, Repeaters, Hubs	Datalink, or MAC* addresses
1	Physical	Unshielded Twisted Pair, Shielded Twisted Pair, Coax, Twinax, Serial cable	Modems, CSU/DSUs	N/A (cables don't have addresses)

*MAC, in this case, stands for Media Access Control, not to be confused with an address for a Macintosh...

Combinations that include a term from each layer describe fully how a packet is getting from a given point "A" to a directly connected point "B". For example, A may be talking to B using IP over Ethernet over Unshielded Twisted Pair; or, "my computer talks to my ISP using IP over PPP over a serial cable" (a modem is simply a serial cable extender in this sense.) From the physical layer standpoint, devices have no addresses. On the datalink layer, all Ethernet and Token Ring cards all have 6-byte addresses manufactured into them, called MAC addresses (nothing to do with

Macintoshes.) Point-to-point links such as serial lines do not have MAC addresses, which creates special cases from a data transmission standpoint, that are outside the scope of this document.

The Physical layer defines the electrical media and signaling used to transmit information on a wire (or wires.) The datalink layer defines the format of the data as it is transmitted (e.g., an Ethernet frame.) Network layer information is encapsulated inside datalink layer frames. If you look at an IP packet on an Ethernet wire it would look something like this:

Ethernet Header (with dest and src MAC addr)	IP Header (with dest and src IP addr, and checksum)	Actual Data
--	---	-------------

Note that this indicates that, in order for two Ethernet-attached stations to communicate with each other via IP, they must know the MAC address of each other. If station "A" knows the IP address of station "B", and knows station "B" is on the same subnet, station "A" will issue an Address Resolution Protocol (ARP) broadcast. An ARP broadcast is a message that says, "Who out there is 192.168.1.1?" The TCP/IP software running on the workstation or router at 192.168.1.1 is responsible for sending back an ARP response that says, "I am 192.168.1.1, and my MAC address is 08:00:09:AF:24:33." All stations keep an ARP cache with the MAC and IP addresses of all the stations it recently communicated with directly. Try the command "arp -a" sometime on a UNIX or Windows workstation; on a Cisco router, the command is "show arp".

Note that layer 1 devices are "invisible" to layer 2; and layer 2 devices are "invisible" to layer 3. In other words, TCP/IP doesn't care if you're running over Ethernet or Token Ring, as long as it's connected properly. In fact, you can put bridging and/or switching devices on your network without disturbing any of your IP subnetting. Similarly, you can convert between different types of media (e.g., coax to twisted pair) without any layer 2 devices being aware of the change. To change layer 1 media, you typically need a layer 2 device (e.g., "I have a Ethernet Coax to Ethernet Twisted-Pair repeater".) To change the layer 2 protocol (e.g., Ethernet to Token Ring) you typically need a layer 3 device (a router.) All this is good, since it allows some measure of media independence within the network; you can run IP over just about anything better than two cans and a string, and even that, if you can find transceivers to handle it ;-)

4. Intro to Ethernet

Developed in the early 1970's, Ethernet has proven to be one of the most simple, reliable, and long-lived networking protocols ever designed. The high speed and simplicity of the protocol has resulted in its widespread use.

Although Ethernet works across a variety of layer one media, the three most popular forms are 10BaseT, 10Base2, and 10BaseF, which use unshielded twisted pair (UTP), coaxial, and fiber optic cables respectively. UTP is used in a "star" configuration, in which all nodes connect to a central hub. 10Base2 uses a single coaxial cable to connect all workstations together in a "bus" configuration, and does not require a hub. 10BaseF uses fiber optics, which, though expensive, can travel long distances (2km) and through electrically noisy areas.

An interesting difference between coaxial Ethernet and other types is that coax Ethernet is truly a one-to-many (or, 'point-to-multipoint') connection; fiber and UTP connections are, from a layer one

perspective, one-to-one (or, 'point-to-point') connections, and require an additional networking device (typically, a repeater, or Ethernet hub) to connect to multiple other workstations. This is why coax Ethernet does not require a hub, and Ethernet over other media typically does.

Ethernet Topologies			
	Pro	Con	Typical Use
10BaseT	*Very reliable- one fault usually doesn't affect entire network.	*Relatively short distance from hub to workstation (100m). *Requires a lot of wiring (a separate link for each workstation.)	*Offices and home networks.
10Base2	*Cheap- no hub required, no wiring except from station to station. *Well shielded against electrical interference. *Can transmit longer distances (200m).	*Any break in connectivity disrupts entire network segment. *Problems can be very difficult to troubleshoot.	*Small or home networks, hub to hub links.
10BaseF	*Long distance networking (2000m). *Immune to electrical interference.	*Very expensive to install.	*Long distance hub-to-hub or switch-to-hub links.

Ethernet is like a bunch of loud people in an unmoderated meeting room. Only one person can talk at a time, because communication consists of standing up and yelling at the top of your lungs. People are allowed to start communicating whenever there is silence in the room. If two people stand up and start yelling at the same time, they wind up garbling each others' attempt at communication, an event known as a "collision." In the event of a collision, the two offending parties sit back down for a semi-random period of time, then one of them stands up and starts yelling again. Because it's unmoderated, the likelihood of collisions occurring increases geometrically as the number of talkers and the amount of stuff they talk about increases. In fact, networks with many workstations are generally considered to be overloaded if the segment utilization exceeds 30-40%. If the collision light on your hubs is lit more often than not, you probably need to segment your network. Consider the purchase of a switch, described below.

Ethernet hubs are used in 10BaseT networks. A standard hub is just a dumb repeater-- anything it hears on one port, it repeats to all of its other ports. Although 10BaseT is usually wired with eight wire jacks (known as RJ45 connectors), only four wires are used-- one pair to transmit data, and another pair to receive data. While transmitting, an Ethernet card will listen to its receive pair to see if it hears anyone else talking at the same time. These two behaviors (listen for silence before talking, and detect other people talking at the same time) are described by the acronym people as CSMA/CD, or "Carrier Sense Multiple Access, Collision Detection."

One hundred megabit Ethernet (100BaseTX) works just like ten megabit Ethernet, only ten times faster. On high-quality copper (known as Category 5, or CAT 5 UTP), 100BaseTX uses the same two pair of copper to communicate. If you have standard network-quality copper, an alternative is to use 100BaseT4, which uses all four pairs, but can communicate at 100Mbps on CAT 3 UTP.

Gigabit Ethernet works just like hundred megabit Ethernet, only ten times faster (1000Mbps, or 1Gbps.) There are some Gigabit Ethernet devices floating around out there, but no standard had been created, so early adopters are likely to find their Gigabit Ethernet devices in need of replacement or upgrade when a standard is ratified.

If your conference room gets too busy, you may consider splitting them into two groups by putting a partition wall with a door between the halves, and putting a person in the doorway. This person would listen to the conversations in both rooms, memorize the names (Ethernet card addresses) of everyone in each room, and forward messages from room to room when necessary. A device to do this is called a "transparent bridge." It's called "transparent" because it's smart enough to learn the Ethernet addresses on its own without the workstations suspecting anything is going on. ["Source-route bridges" are uncommonly used so I'm not going to discuss them.]

Ethernet switches are little more than high-speed, multi-port bridges. They learn the Ethernet addresses of everyone attached to each port, and make intelligent forwarding decisions based on Ethernet card address (aka MAC address.) Because communication between 100Mbps and 10Mbps networks requires buffering, Ethernet switches are often used for this purpose. Many inexpensive switches have many 10Mbps ports and one or two 100Mbps ports. Typically, you would connect your server(s) to the 100Mbps port(s), and workstations or entire hubs to the 10Mbps ports. The buffering and intelligent forwarding allows another interesting feature to exist-- "full-duplex" Ethernet. "Half-duplex" means you can either talk or listen, but not both, at a given time, such as when using a radio. "Full-duplex" communication means you can talk and listen at the same time, such as when on the phone. Since 10BaseT uses separate pairs of copper for sending and receiving, it's physically possible to do both if there are no other workstations on your network segment-- which is the case if you are directly attached to a switch. Note that both the switch port and your network card must be configured for full duplex operation for this to work, but the result is worth it: a full 20Mbps for "regular" Ethernet and a whopping 200Mbps of bandwidth available for full-duplex fast Ethernet. Since collisions are eliminated, the 30% rule does not apply. When considering the purchase of a switch, there are a few important considerations, not all of which may apply to your requirements:

- Does the switch support 100Mbps on any ports? How many, and will it autodetect 10/100BaseT?
- Does the switch support full duplex? Even on the 100Mbps ports?
- How many MAC (Ethernet card) addresses does it store? 500? 5000?
- Some "workgroup" switches only allow one MAC address per port, so these would not be suitable if you plan to connect hubs to switch ports.
- Some hubs are advertised as "switching hubs" but actually require the purchase of a separate "switch module" to function as such. I won't name names, but I'll never buy a hub, nor anything else, from that manufacturer, if given a choice. I consider it to be an outright lie.
- You tend to get what you pay for. If a switch seems unreasonably inexpensive compared to other switches that appear to have similar specs, look closer, or check the detailed specs on the manufacturer's web site. Often, you'll find that a cheap switch either isn't a switch at all (see last item) or only allows one workstation per port (see item above last item.)

<SOAPBOX>

I have increasingly seen people install 100Mbps networks without paying any attention to whether or not there is a need to do so. Most smallish networks do not need 100Mbps switched Ethernet; in many cases, excellent results can be obtained by purchasing a 10/100 Ethernet switch. Connect the

100Mbps port to the server, and the 10Mbps ports to the workstations or hubs. You greatly increase the amount of bandwidth available, without pulling new cable and installing new cards in the workstations. Even switched, full-duplex 10Mbps Ethernet increases the available bandwidth by almost 600%. People talk about the panacea of "reduced latency" that switched networks provide, but most modern protocol implementations are designed to be almost completely unaffected by a few milliseconds of latency. Most computing environments are disk I/O bound, not network or CPU bound; yet people will recable their networks and install new network cards, or buy servers with faster and faster (read: more and more idle) processors, but the most performance benefit typically lies in installation of more memory in the server, or addition of a caching RAID controller and a RAID-based disk subsystem. Before blindly "upgrading stuff" to improve the speed of your network, try to find out where the true bottleneck lies.

</SOAPBOX>

5. Why is IP so much more difficult than IPX?

I have gotten some interesting feedback on the title of this section. From a LAN administrator's standpoint, IPX is almost completely auto-configuring. Since TCP/IP requires substantially more administrator understanding and time to properly implement, then IP, from a LAN administrator's standpoint (this document's target audience), is substantially more difficult to work with than IPX. You don't find 15+ page documents on the Internet about "the fundamentals of IPX", do you?

The four items you need to use IP effectively on the Internet (that you don't need to set up an IPX workstation) are the IP Address, the IP Subnet Mask, the IP Address of the Default Router, and the IP Address(es) of your Domain Name Servers (DNS Servers, often shortened to "Name Servers.")

IP Addresses: IP uses 4-byte addresses, like 192.168.1.1. IPX uses 10-byte addresses, like 10000001:0000C04C1141. Those happen to be the IP and IPX addresses of the workstation I'm using now. "But wait," you ask, "I've used IPX before and all it uses are four byte addresses." Well, that's not entirely correct. The 4-byte "IPX Address" configured into IPX-based servers is only the *network* portion of the address. All addresses used by routable protocols have a "network" portion, which gets your packet to your nearest router, and a "host" portion, which indicates which host station you are on that routed segment. The 4-byte "IPX Address" you define is actually a 4-byte "IPX *Network* Address." The other 6 bytes is the hardware address of your NIC. Since IP addresses don't use the unique hardware address of your NIC, you must define them manually (or semi-manually by configuring a BOOTP or DHCP server, a task which is currently outside the scope of this document.)

IP Subnet Masks: Subnet masks (described in more detail in the next section) are used in IP to determine which part of the four-byte IP address describes the network you're on, and which part describes which host you are on that network segment. In IPX, the first four bytes always indicate the network you're on, and your six byte MAC layer address indicates which host you are on the network segment. In IP, the portions used to describe which network you're on can range from the first 8 bits of the address, to including all except the last two bits of the whole address. More in the next section.

Default Router: In IPX, routers are identified by sending out a broadcast that says, in essence, "Hey? Who out here is a router?" In IP, there has historically NOT been any automatic method for router discovery. There is now a protocol for IP router discovery, but it is not widely implemented. Therefore, you must tell the workstation what the address of the local router is. Note that with end-station PPP (like Win95 Dial-Up Networking), the default route is automatically set to, "out the serial

cable." You do not need to set more than one default route. If the default router feels the packet would reach a destination better through a different router, the default router will tell your IP stack to use the other router (this is an ICMP Redirect.) If you specify no default route, no packets from that workstation can make it off the local wire; therefore, it is better to set a *wrong* default route than *no* default route. If in doubt, set the default route to the address of any known router on the local subnet.

DNS: In IPX, designed by Novell, the names (and corresponding addresses) of ALL services available on the network are stored in ALL Netware servers as a SAP table (SAP stands for Service Advertising Protocol.) Netware servers will share SAP information with each other automatically. Unfortunately, since ALL servers must know about ALL services, SAP tables can get very unwieldy on large networks, and without the benefit of advanced routing/advertising algorithms (NLSP), can flood networks with SAP broadcasts. The way IP handles name-to-address translation is called DNS. When you query your DNS server for a given name's address (such as www.novell.com), the DNS server will query one of the "root" servers for .COM. The root server tells the DNS server the address of the "authoritative" DNS server for novell.com. Your DNS server then asks the DNS server of novell.com what the address of www.novell.com is; when novell.com's DNS ponies up the address of www.novell.com, your local DNS "remembers" where www.novell.com was, so it doesn't have to look again the next time someone asks for that name's address. Note that DNS uses special records for mail routing, called MX records, that usually differ from the host addresses. Therefore, an ftp (or www, or gopher,...) connection to microsoft.com probably reaches a different address than mail sent to somebody@microsoft.com. Of course, the giveaway that you're talking mail ("MX" record) addresses, rather than host ("A" record) addresses, is the "@" in the address. Host names never have @ symbols, which is why you connect to www.microsoft.com, never www@microsoft.com.

BOOTP and DHCP: BOOTP was designed to ease the configuration of desktop IP stacks. In a nutshell, a BOOTP-enabled workstation sends out a broadcast BOOTP request, which is answered by a BOOTP server. The answer includes workstation address, subnet mask, default route, and DNS location(s). DHCP is generally accepted as the "next generation" of BOOTP. Whereas BOOTP statically assigns IP addresses by MAC address, DHCP supports address "leasing" where an address is granted to a specific MAC address for a finite amount of time, and can be reused after a specified amount of time. DHCP also supports fields beyond BOOTP, most notably returning information about the location of WINS server to Windows NT clients, and the location of DSS servers to Netware/IP clients. (A DHCP service is included with NT, and is available for free download as part of the Netware/IP upgrade for Netware 4.10 servers, see <http://support.novell.com>.)

6. IP Addresses, Subnet Masks, and Subnetting

Part A: The World According to RFC 950 (the current/old way of doing things)

An IP Address is broken up into three parts: the network portion, the subnet portion (optional), and the host portion. The size of the network portion is determined by the first byte of the address:

First Byte	Class	Network Mask (explained later)
1-126	"A"	255.0.0.0
128-191	"B"	255.255.0.0
192-223	"C"	255.255.255.0

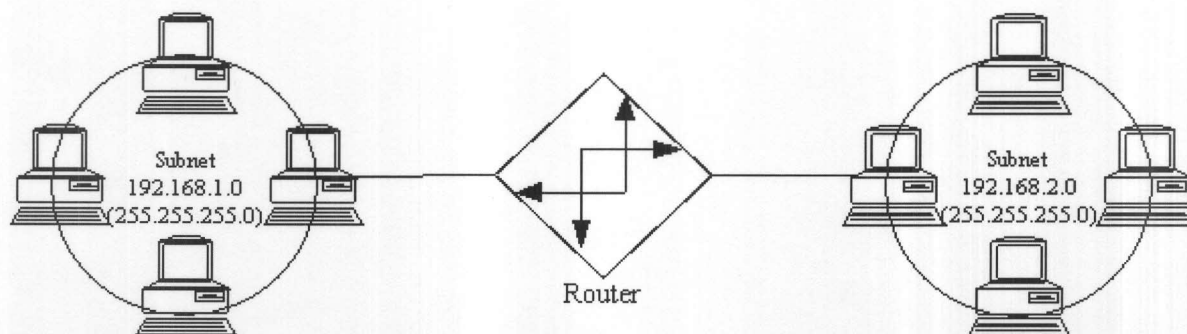
Note: people often refer to any subnet with a mask of 255.255.255.0 as being a class "C" network; however, the only "true" class "C" networks have a first byte in the range of 192-223. This becomes important when you start subnetting.

The Subnet portion of an IP address is actually optional, and, in fact, is rarely used on class "C" networks. Generally, you can subnet any network you have control over, in any valid way you want. The tricky part is understanding what is valid.

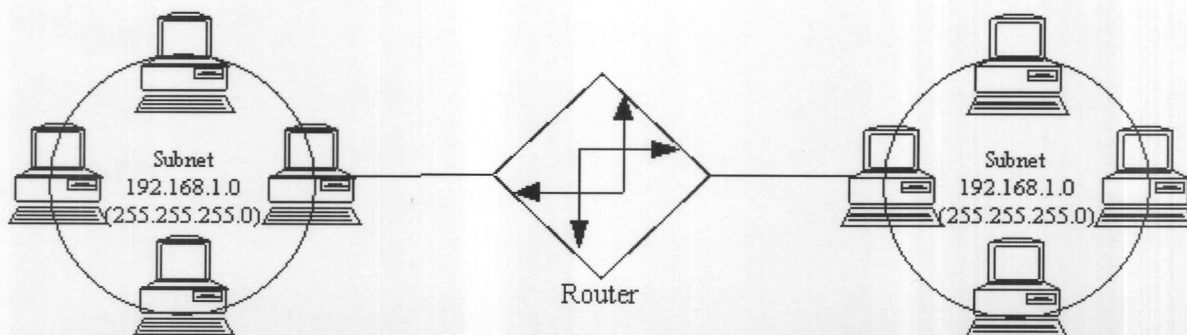
Lets start with some ground rules:

- All hosts on the same subnet must agree on the subnet mask, particularly the routers. Otherwise, packets actually intended for another subnet may never leave the existing subnet: a host won't give to the router a packet it thinks is destined for the local segment. This behavior is important to understand: the router doesn't automatically forward packets, the hosts have to actually *give* the packets *to* the router.
- No two different subnets can include the same host address. This can get tricky when subnetting in an unusual manner.
- The top and bottom host numbers are reserved; the bottom one (usually ????.0) is shorthand for the whole subnet, and the top one (usually ????.255) is the broadcast address. Some implementations also use .0 as a broadcast address, so it is never safe to use for a host.
- The bits in the subnet portion cannot be all ones. This requires a bit of binary arithmetic to determine which subnets would be invalid

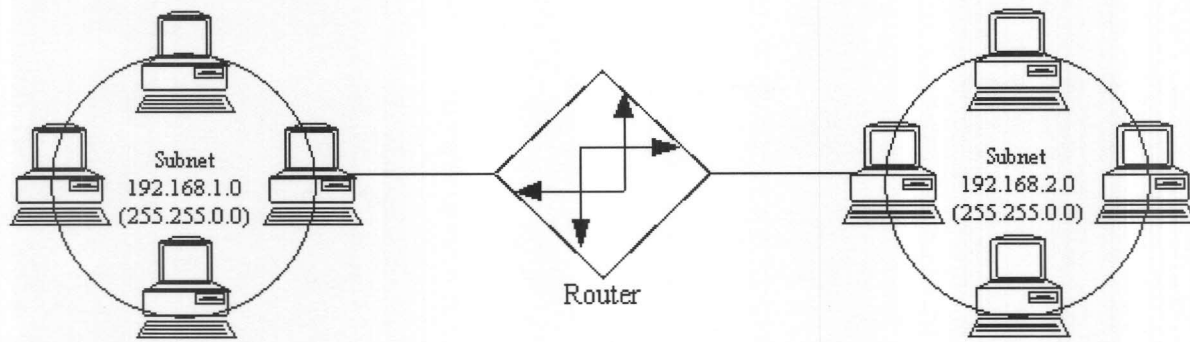
Valid Configuration:



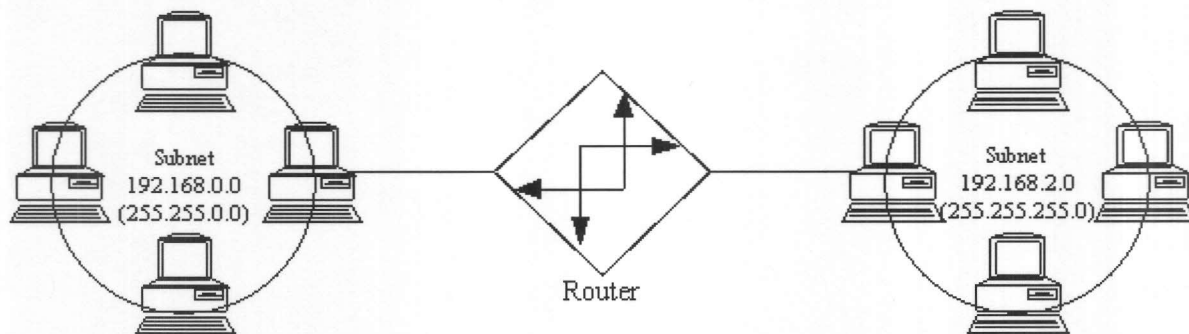
Invalid Configurations:



...This is invalid since the [exact] same subnet exists on both sides of the router.



...This is invalid since the same subnet exists on both sides of the router. Watch that subnet mask!
(See below.)



These images created using SmartDraw. [Click Here](#) for a free trial copy.

...This is invalid because a the same host address could be "valid" on either subnet, e.g. 192.168.2.100. Even though the right side subnet is valid by itself, it is actually a small piece of the left side network. Address overlap is never allowed (which subnet would the router forward a packet destined for 192.168.2.40 to? Both directions are equally valid.)

The Glossy Explanation

When using a subnet mask of 255.255.0.0, the first two bytes indicate the network you're on, and the last *two* bytes indicate the host you are on that network. Very rarely will you find a network segment with 65,534 hosts on it, though. You'll only find network masking like that used closer to the Internet backbone, in the context of, "All them hosts [and subnets thereof] are thataway." Now, that brings up one of the nice features of subnet masking: you can lump a bunch of networks together by using unusual subnet masking; however, that sort of activity generally doesn't happen on the near side of the 'net.

When using a subnet mask of 255.255.255.0, the first three bytes indicate the network you're on, and the last byte is the host you are on that network. Hosts .1 through .254 are available.

By using a subnet mask of 255.255.255.128, you can split that network into two halves, the first half containing the host addresses .1 through .126, the second half containing the host addresses .129 through .254. Note that on a true class "C" network, you can't use the top subnet, since the bit in the subnet portion (one bit on a class "C") would be one (refer to ground rule "D".)

By using a subnet mask of 255.255.255.192, you can split the network into four portions, each with 64 hosts (62 usable.) Subnetwork one includes the addresses .1 through .62, subnetwork two includes the addresses .65 through .126, subnetwork three includes .129 through .190, and subnetwork four includes the hosts .193 through .254. On a true class "C" network, subnetwork four is not valid.

You can not arbitrarily cut a piece out of one network and place it on another segment; the best you can do with a given subnet (or network) is chop it in halves, or quarters, or eighths, or sixteenths... (note the "powers of two" progression; this is an effect of stealing bit positions from the host address section, and giving those bits positions to the subnet portions. It gets complicated...)

Part B: The World According to RFC 1812 (the new way of doing things)

or, By The Way - Forget Everything You Just Learned

Under RFC 1812, things have changed..!

Perhaps the most significant change on the near side of the 'net under RFC 1812 is Classless Inter-Domain Routing (CIDR, pronounced "Cider"). Under CIDR, the concept of separate "network" and "subnet" portions is now considered outdated, and is being replaced by a "classless" addressing scheme where addresses can be "subnetted" more freely, without consideration of the "class" of address. With the removal of the subnet portion, and the liberalization of (what is now called) the network prefix, there is no longer a consideration of whether or not the bits within the subnet portion are all ones; in other words, you no longer lose a subnet when you break up what used to be known as a class "C" network. You can also aggregate formerly class "C" networks together using network prefixes fewer than 24 bits long. For example, you could combine the formerly class "C" networks 192.168.2.0 and 192.168.3.0 into a single subnet with 510 usable addresses, by using a network mask of 255.255.254.0. What you're really saying here is that the last bit of the third byte now belongs to the "host number" portion of the address, and the "network prefix" is 23 bits (two bytes and seven bits) long. Therefore, the two networks being combined must be contiguous, and the third byte must be even on the lower numbered network. You could *not* combine, for example, 192.168.2.0 and 192.168.5.0; not could you combine 192.168.11.0 and 192.168.12.0. You could follow similar rules to combine four contiguous class "C" style networks, but the third byte of the lowest numbered network would have to be a multiple of four. This sort of thing is routinely done (on an increasingly larger scale) as you get closer to the Internet backbones.

Most of the other effects of RFC 1812 and CIDR routing affect areas of the 'net closer to the backbone, and mostly work to reduce the size (or at least the rate of growth) of routing tables in backbone routers.

Part C: Huh? (or, Perhaps you could apply an analogy to all this?)

A good analogy for IP addressing and packet forwarding (routing) is the snail mail analogy. Consider an IP packet to be an envelope containing data, and having an address on the front. Every TCP/IP-enabled network interface can be compared to a mailbox. Every mailbox (interface) has an IP address. The four bytes of an IP address can be compared to the state, city, street, and house number fields on the front of a snail mail envelope. A router in this analogy is a post office, that sorts and forwards mail based on the address on the envelope (packet header.) If the address is on the same street (based on the subnet mask,) the envelope (packet) is sent directly to the destination mailbox

(interface) via local courier (Ethernet?). If the address is determined to be on another street, or in another city or state, the envelope (packet) is delivered via local courier (Ethernet?) to the street's post office (router), where the postal workers (routing software) sort and forward mail based on established post office sorting procedures (routing tables.) The breakdown in this analogy, of course, is that no routing software has ever been known to shoot people. (Just Kidding :-)

7. Subnetting, Bit by Bit

A. Binary arithmetic

You may have heard that computers represent all numbers as "bits", or "zeros and ones." It would be more fair to say that computers work primarily with groups of eight 0's or 1's, called bytes. In practice, most desktop PC's work with clumps of four bytes at a time, or 32 bits. That's why 80386 through Pentium II processors are called 32-bit processors. [Although Pentium class processors have some 64-bit attributes such as a 64-bit external memory bus, they still do most operations as 32-bit operations.]

Now, think back to first grade math, when the teacher was describing the decimal numbering system. As it happens, it's called "decimal" because it's a numbering system that uses ten numbers: the numbers zero through nine. If you need to represent a number larger than nine, you have to start adding digits; then the teacher described the ones place, the tens place, the hundreds place, etc. For example, the number 45678 has a four in the "ten thousands" place, a five in the "thousands" place, a six in the "hundreds" place, a seven in the "tens" place, and a 8 in the "ones" place:

Ten Thousands	Thousands	Hundreds	Tens	Ones
4	5	6	7	8

Since computers work in binary, and only have "0" and "1" to work with, they have to start new digits ("binary places", not "decimal places") as soon as they get past the number one! In decimal, the "decimal places" were all powers of ten:

$10^0=1$,
 $10^1=10$,
 $10^2=100$,
 $10^3=1000$, etc.

In binary, the "binary places" follow powers of two:

$2^0=1$ (1 binary),
 $2^1=2$ (2 binary),
 $2^2=4$ (4 binary),
 $2^3=8$ (8 binary),
 $2^4=16$ (16 binary),
 $2^5=32$ (32 binary),
 $2^6=64$ (64 binary),
 $2^7=128$ (128 binary),
 $2^8=256$ (256 binary), etc.

The number 45678 is represented in binary as follows:

(Binary Places, expresses as Decimal:)	32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1
	1	0	1	1	0	0	1	0	0	1	1	0	1	1	1	0

(Add up the columns where you find ones: 32768 plus 8192 plus 4096 plus 512 plus 64 plus 32 plus 8 plus 4 plus 2 equals 45678!)

Counting to Forty:

Decimal	Binary	Decimal	Binary	Decimal	Binary	Decimal	Binary
1	1	11	1011	21	10101	31	11111
2	10	12	1100	22	10110	32	100000
3	11	13	1101	23	10111	33	100001
4	100	14	1110	24	11000	34	100010
5	101	15	1111	25	11001	35	100011
6	110	16	10000	26	11010	36	100100
7	111	17	10001	27	11011	37	100101
8	1000	18	10010	28	11100	38	100110
9	1001	19	10011	29	11101	39	100111
10	1010	20	10100	30	11110	40	101000

Now, an IP Address is four bytes, eight bits each, represented as decimal numbers with periods in between; for example, 10.5.72.230. This number can be represented in binary (remember when I said that IP Addresses are best expressed as 32-bit binary numbers? I did mention that, didn't I?) as b00001010.00000101.01001000.11100110. (The "b" means "binary"; that and the periods are added for your convenience.) Now, 2^{32} (two to the thirty-second power) is 4294967296, or just over four billion. So, theoretically, there are over four billion IP addresses available to the world; so why is there a shortage? (Oh yeah, have you heard? There's a shortage. Last I checked, they're projecting to run out of IP addresses around the year 2025.) Well, as it turns out, trying to keep track of where four billion individual hosts are would be pretty much impossible for equipment today, and certainly impossible for equipment ten years ago when this was being developed. So, routing was (over) simplified by splitting the IP address space into "classes"; those IP addresses whose first byte was in the range 1-126 would belong to networks of 16777214 ($2^{24}-2$) hosts; these were called "Class A" networks, and there are 127 of them. In Class A networks, the first eight bits are the "network portion", and the last 24 bits are the "host portion." Those IP addresses whose first byte was in the range 128-191 were called "Class B" networks of 65534 ($2^{16}-2$) hosts, and there were 16384 (that's $(192-128)*256$) of them. That's 16 bits for the network portion, and 16 bits for the host portion. "Class C" networks, where the first byte is in the range 192-223, have a 24 bit network portion, and an 8 bit host portion. Note how neatly everything lines up on byte boundaries:

Class	Network bits	Network Mask	Network Mask (binary)
A	8	255.0.0.0	b11111111.00000000.00000000.00000000
B	16	255.255.0.0	b11111111.11111111.00000000.00000000
C	24	255.255.255.0	b11111111.11111111.11111111.00000000

Now, since it's unlikely that a network administrator is going to want to have some 16777214 (nearly seventeen million) hosts on the same network segment(!), network administrators were allowed to administratively split up their networks by subnetting them. Routing on the Internet backbones was fairly simple... until they started to hit the Class C networks hard. If your company needed 1000 IP addresses, you'd probably get four Class C networks to accommodate them... but that would add four individual routes propagated to every "backbone" router on the Internet! Hence the need to split up networks on other than just byte boundaries.

This is where everything got hard.

It turns out that you can combine four "Class C" networks together into one routing table entry by using a subnet mask (aka Network Prefix) of 255.255.252.0. But not just any four; as it happens, they must be contiguous, and the third byte of the first network must be a multiple of four (like the number 204 is.) If you want to join eight of them together, the first network must be a multiple of eight (which the number 204 is not.) If you want to join ten networks together... well, you can't. Ten is not a power of two. Funny how everything follows powers of two...

B. Boolean Logic and The Binary "AND"

Named after the nineteenth-century mathematician George Boole, Boolean logic is a form of algebra in which all values are reduced to either TRUE (1) or FALSE (0). All math performed by modern computers is done using Boolean algebra. A few basic operations:

Operation	Result	Examples
AND	true if A <i>AND</i> B are true	1 AND 1 = 1 1 AND 0 = 0 0 AND 1 = 0 0 AND 0 = 0
OR	true if A <i>OR</i> B are true	1 AND 1 = 1 1 AND 0 = 1 0 AND 1 = 1 0 AND 0 = 0
XOR (eXclusive Or)	true if <i>either</i> A <i>or</i> B are true	1 XOR 1 = 0 1 XOR 0 = 1 0 XOR 1 = 1 0 XOR 0 = 0
NOT	opposite of A	NOT 1 = 0 NOT 0 = 1

The binary "and" operation is often used when you want to see only certain bits of a given byte-- a procedure called "masking." Some of you may have seen a similar thing in school; some of my teachers used to conduct multiple-choice tests where you would fill in a circle corresponding to the answer I thought was correct. The teacher would then take an overlay, or mask, and place it over the answer sheet. This overlay had holes only where the marking spots for the correct answers were, and the teacher would mark any answers where he/she didn't see a mark as incorrect. The subnet mask is used in this fashion by the computer to determine which bits are the network portion of an IP address, and which bits are used for the host, or workstation, portion.

C. The Subnet "Mask"

The subnet mask is used to figure out what network you're on. The reason it's called a "mask" is the same reason the tape you use to cover trim when painting is called "masking tape"; you use it to cover up the parts you don't want to deal with right now. Did you notice how, in a binary AND, any time B is zero, the result is zero? And any time B is one, the result is whatever A is? Hmmm.....

The primary use of the subnet mask (from the perspective of the Near Side of the 'Net) is for workstations to determine whether or not the server or workstation they're trying to talk to (the "destination IP address") is on the same subnet as itself; if they destination IP address is on your subnet, you'll send the IP packet directly to the other computer via the Ethernet or Token Ring (or whatever) network you're on, without bothering the router... at all! **The first routing decision made on an IP packet is made by the workstation sending it; it decides whether or not to send the packet to a router.** Doing this is a four step process:

1. **Step 1:** Convert the IP Addresses to Binary.
If necessary, the IP address is converted from the familiar dotted-decimal into a 32-bit binary value. It sucks as much for the computer to do it as it does for humans to do it, but computers generally complain less, and they're good at math :-)
2. **Step 2:** Apply Source subnet mask to Source addresses:
The **network portion** of the **workstation's IP address** is determined by performing a binary AND operation on the workstation's IP address and its subnet mask. This operation "masks off" all of the bits of the "host portion" of the IP address, and leaves the "network portion" behind for comparison with the destination's network portion. Hey, wait a minute? How do we know what the subnet mask of the destination is?
3. **Step 3:** Apply Source subnet mask to **Destination** addresses:
As it happens, we don't care what the subnet mask of the destination is. We only care if the destination is on our same network segment! **Since every workstation on our network segment shares the same subnet mask, we can apply our subnet mask to the destination to determine if its network portion matches ours.** So, the **network portion of the destination workstation's IP address that we can use to see if it matches ours** is determined by performing a binary AND operation on the destination IP address and **our subnet mask**.
4. **Step 4:** Compare the derived **network** portions for equality:
At this point, we can compare the network portions we have masked from the source and destination IP addresses to see if they're the same. If they are, then we **must be on the same subnet** so we send the packet directly; if they are different, even by only one bit, the destination is on another network segment...somewhere. We don't know where. Maybe the router does...

OK, so let's try this a few times ourselves; get a few IP addresses and subnet masks together and plug 'em into Daryl's Subnet Calculator! Requires JavaScript to be enabled on your browser. If you're

reading a hard copy of this, the full URL is <http://ipprimer.windsorcs.com/subnet.html>.

Remember the part about combining four "Class C" networks together? Watch your binary arithmetic:

(network prefix bits shown in green)

Networks	Networks, in Binary
192.168.8.0	b11000000.10101000.00001000.00000000
192.168.9.0	b11000000.10101000.00001001.00000000
192.168.10.0	b11000000.10101000.00001010.00000000
192.168.11.0	b11000000.10101000.00001011.00000000
Mask, 255.255.252.0	b11111111.11111111.11111100.00000000

Notice how all of the bits above the ones in the subnet mask stay the same; following the rules above, all hosts on these networks, if you apply the mask, are on the same network. This was called "supernetting", but now is called "CIDR Routing", pronounced "Cider Routing".

Doing it wrong:

(network prefix bits shown in red)

Networks	Networks, in Binary
192.168.10.0	b11000000.10101000.00001010.00000000
192.168.11.0	b11000000.10101000.00001011.00000000
192.168.12.0	b11000000.10101000.00001100.00000000
192.168.13.0	b11000000.10101000.00001101.00000000
Mask, 255.255.252.0	b11111111.11111111.11111100.00000000

Oops-- seems the sixth bit of the third byte changed within the network prefix portion (the part above the 1's in the subnet mask), so with the given subnet mask (255.255.255.252), 10.0 and 11.0 would ALWAYS be on a different network aggregation than networks 12.0 and 13.0. Confused? Play with it in the [Subnet Calculator](#), and compare the network portions.

D. "Slash" Notation

Subnet masks are often abbreviated using a forward slash "/" and the number of ones in the mask. For example, a network 192.168.1.0 with a subnet mask of 255.255.255.0 can be expressed as 192.168.1.0/24 (since 255.255.255.0 is 24 ones followed by eight zeros.) Therefore, a /25 subnet is a subnet with a mask of 255.255.255.128, and a /26 subnet has a mask of 255.255.255.192, etc.

E. A Neat Trick

Now that you actually understand the binary arithmetic behind subnet masking (well, I hope you do, anyway) we can cover some of the neat tricks for computing subnet masks. To determine the number of hosts on a given subnet (assuming the subnet is smaller than class "C",) simply subtract the last number of the subnet mask from 256. For example, a subnet mask of 255.255.255.224 has 32 hosts

($256-224=32$.) Then you can just divide the result into 256 to determine the number of subnets ($256/32=8$.) So, using a subnet mask of 255.255.255.224 gives you 8 subnets of 32 hosts each. Of course, this only works when you are subtracting a number that is a power of two (1, 2, 4, 8, 16, 32, 64, or 128.) When the network prefix is larger than class "C", you can determine how many class "C" networks are aggregated by subtracting the third byte from 256-- so a network prefix of 255.255.240.0 is an aggregation of ($256-240$) 16 class "C" networks.

Thanks to Gael M. for this tip.

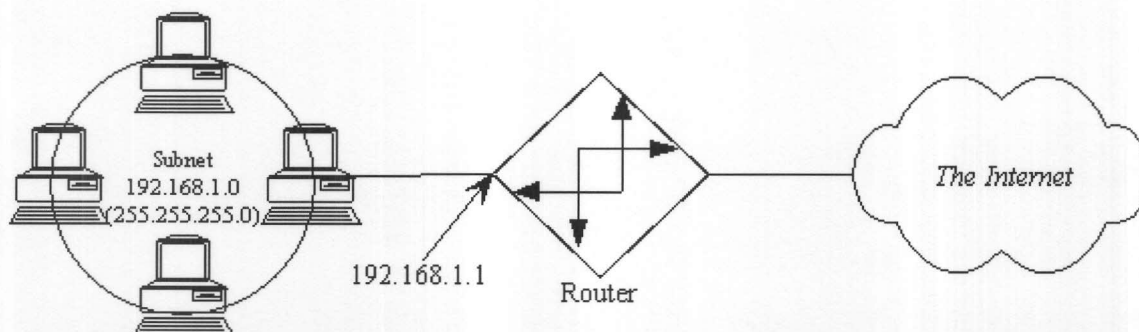
F. In closing...

Why all this crap about binary arithmetic? Do I *have* to know this stuff? I'm afraid so; subnet masks are created and used on a bit-by-bit basis; in order to effectively use subnet masks that *don't* fall on byte boundaries (like 255.255.255.0 does), you have to determine what hosts are on each subnet by using binary arithmetic. It sucks, it's hard, it's confusing (especially since IP addresses and masks are expressed in decimal instead of hexadecimal notation) but you must use and understand IP addresses and subnet masks as binary.

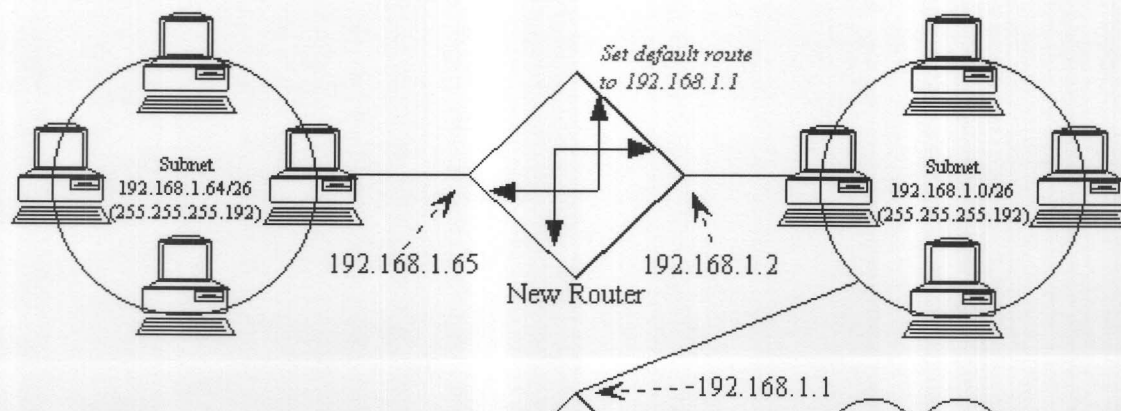
8. Routing and Static Routes

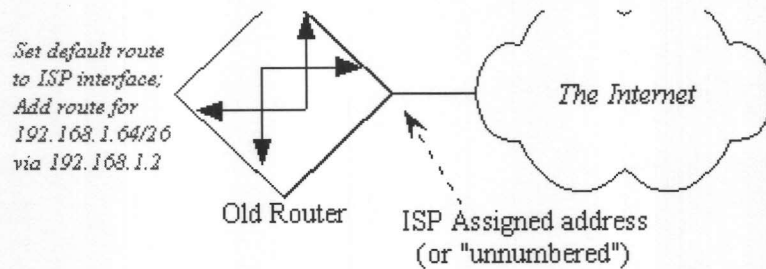
I'm not going to go into a ton of detail here. Instead, I'm going to offer a single example of a network split into two halves.

Before: Network 192.168.1.0:



After: Split into three parts using a subnet mask of 255.255.255.192





These images created using SmartDraw. [Click Here](#) for a free trial copy.

What we need to do now is tell the router what happened...

First, you have to tell the old router that the network attached to its Ethernet interface has changed (specifically, the network mask has changed, and often, the address of the Ethernet interface has changed.) If you were adding a new subnet, rather than splitting an existing one, then you could probably skip this step.

Second, you have to tell the old router where to find the new network (what the next hop is.) A typical command would look something like this:

```
ROUTE 192.168.1.64/255.255.255.192 192.168.1.2
```

What you're telling the old router with that statement is, "if you need to route packets to the subnetwork that starts at 192.168.1.64 and has a subnet mask of 255.255.255.192, you should forward all packets intended for that network to the router at 192.168.1.2."

Third, be sure the default route for the new router is set to 192.168.1.1.

Note that the automatic routing protocol (IP) RIP does not understand subnet masking. If you are using protocols that do, such as OSPF or EIGRP, then you probably aren't reading this document. *Actually using* routing protocols tends to be irrelevant on the "near side" of the net, since there is generally only one path to the Internet from any given workstation on a LAN. Multiple routes tend to be a problem only closer to the backbone, and that's your ISP's problem.

9. Troubleshooting

The most useful tool in troubleshooting client IP issues is PING. Ping is a low-level method of determining if a specific host is alive.

Step #1: Determine if the IP stack is alive. There is a reserved address 127.0.0.1 called "localhost". A successful ping to 127.0.0.1 means your IP stack is working properly. A ping to localhost doesn't even make it on the wire.

Step #2: Determine if you can talk onto the wire. Ping yourself. If your address is 192.168.1.1, then ping 192.168.1.1. Actually, the packet may or may not actually make it on the wire, depending on your implementation. But it doesn't hurt.

Step #3: See if you can ping anyone else. Ping your default router. Make sure your default router is

on your same subnet! The easy way to do this is to refer to the "glossy explanation" of subnetting in Section 4, and to make sure both addresses *can* exist in the same subnet. If you can't ping your default router, either the router is down (easily checked from another workstation) or there's something wrong at your workstation. Make sure your workstation has the subnet mask set correctly, and that you and the router are using the same frame type. The default frame type for TCP/IP is Ethernet_II on Ethernet LANs, and TOKEN-RING_SNAP on Token-Ring LANs. Cisco routers refer to Ethernet_II as encapsulation type ARPA.

Step #4: See if you ping the far interface of the default router. All routers have more than one interface (or they wouldn't be routers, right?) If you know the interface of the far side of the router, ping that. That verifies that your default route is set properly. If you don't know the address of another router interface, skip to step 5.

Step #5: Ping the address of you name server. Your name server address is given to you by your ISP. If you cannot ping your name server, try to trace your route to it. The UNIX version of the command is "traceroute". The Win95/WinNT version is called "tracert". An example:

```
D:\WINDOWS>tracert ns.orbis.net

Tracing route to ns.orbis.net [205.164.72.2]
over a maximum of 30 hops:

 1  1 ms  1 ms  1 ms  192.168.1.254
 2  60 ms  61 ms  64 ms  205.164.75.1
 3  64 ms  62 ms  65 ms  tamino.summit-ops.orbis.net [205.164.72.129]
 4  78 ms  77 ms  78 ms  ns.orbis.net [205.164.72.2]

Trace complete.

D:\WINDOWS>
```

Note: if you actually get names, you not only have verified Internet connectivity, but you also know your DNS is properly set up. Congratulations! You are on the Internet. If you have problems at this point, it's time to call your ISP.

Step #6: If you didn't get any names in your route trace, don't panic: Try to ping www.novell.com or www.microsoft.com. If you can ping, by name, either of those addresses, you are set up for Internet access. If you get a message like, "Unable to resolve novell.com" then you need to make sure your DNS is set up properly. If you get a "host unreachable" then you probably are set up OK but the 'net

is just a bit congested. (Or you haven't set your workstation's default route properly.)

Typically, I start with step #6, and if that fails, go to step #1.

10. TCP and UDP Communication

TCP and UDP are layer 4 protocols that help organize process-to-process communication. When a Web browser establishes a connection to download an HTML document from `www.mydomain.com`, the browser

1. Resolves the IP address for `www.mydomain.com`
2. Opens a TCP connection to port 80 on the web server `www.mydomain.com`
3. Transfers the data over the TCP connection
4. Closes the TCP connection

Every TCP (or UDP) communication has a source port and destination port number in the TCP (or UDP) header. **Every TCP/IP communication can be uniquely identified as [Source IP]:[Source Port] <--> [Dest. IP]:[Dest Port].** This is how a Web browser can load several images at once and keep track of which packet is for which image. The source port is different for each TCP image-download connection, though the destination port is 80 in each case. For example:

Source IP	Source Port	Dest IP	Dest Port	Notes
192.168.1.1	1025	10.101.10.1	80	index.html
192.168.1.1	1026	10.101.10.1	80	logo.gif
192.168.1.1	1027	10.101.10.1	80	backgrnd.gif

Note that each file getting downloaded has a different source port number; this is how the communications are differentiated (this packet is part of `logo.gif`, this packet is part of `index.html`, etc). Now, let's assume that `index.html` is finished, but the graphics are loading slowly. While the user is waiting, he/she decides to open a telnet session to `rs.internic.net`. The table of open sessions would look like this:

Source IP	Source Port	Dest IP	Dest Port	Notes
192.168.1.1	1026	10.101.10.1	80	logo.gif
192.168.1.1	1027	10.101.10.1	80	backgrnd.gif
192.168.1.1	1028	198.41.0.9	23	telnet rs.internic.net

Now, I could go into exhaustive detail on how a TCP connection is set up and torn down, flow is controlled, and dropped packets are resent. Instead, I'll just say that TCP connections are set up and torn down, and there is flow control and automatic dropped packet redelivery. TCP is like certified mail; if no return receipt is gotten, the packet is resent (I'm oversimplifying but it's close enough.) TCP is used for "reliable" communications, where all data must get through, and must get there in the correct order.

A UDP packet, on the other hand, is more like junk mail. No effort is expended to make sure it

arrives at the destination, or that all packets arrived that were sent. UDP is generally used for real-time applications like Internet radio and online gaming, where dropped packets need not be resent, and would probably be old if they were. UDP is also used when upper-layer protocols do their own flow control and data stream checking and correcting, as is the case in NCP/IP (Netware/IP) and SMB/IP (Microsoft Networking).

Web, Telnet, Mail, and other servers "listen" for new communications at "well-known" TCP port numbers. A short list:

Service	"Well-Known" Port Number
FTP	21&20 (don't ask)
Telnet	23
SMTP Mail	25
HTTP (Web)	80
POP3 Mail	110
News	119
IRC	6667

A more complete list of assigned Well-Known Ports can be found at <http://www.con.wesleyan.edu/~triemer/network/docservs.html>

Publicly available services are generally *always* reached by connecting to their well-known port numbers.

11. Network Address Translation (NAT)

Network Address Translation, or NAT, is accomplished using software that can hide one or more subnets behind a single IP address. NAT software is typically found in newer Internet routers and almost always used in firewalls and proxy servers. NAT is not the same as an HTTP Proxy server. HTTP Proxy servers must be configured on the client side. Once configured, your Web browser asks the HTTP proxy to make connections to the Internet on your behalf; as far as the Web site you're connecting to knows, it's the proxy server that's reading the Web page, not your browser. NAT *is* an effect of the HTTP proxy in this case; the requests from all of the browsers using the HTTP proxy appear to be coming from the proxy server, not from the workstation. The workstation does not need to be using IP addresses that are routable to the Internet; in fact, it is normal to use addresses that are reserved for this purpose, such as 10.x.x.x (*see* Tips and Tricks, later in this document.)

"Transparent" NAT is easier to implement (since nothing needs to be changed at the workstations). However, "configured" NAT (e.g., HTTP proxy servers) often add additional features, such as Web page caching.

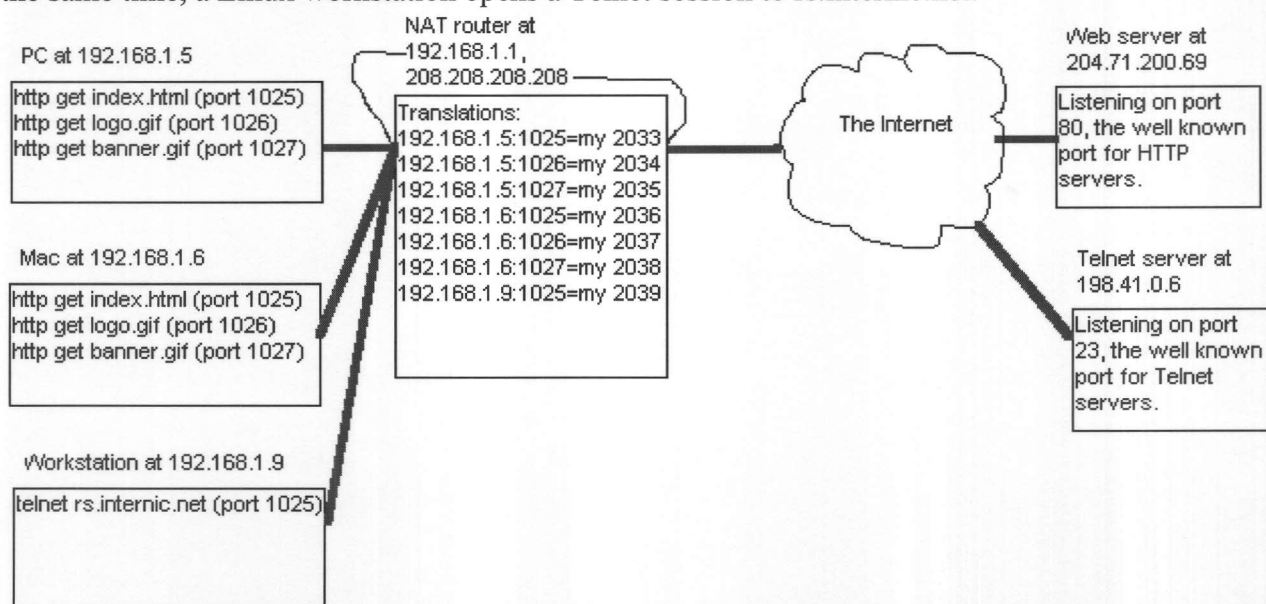
NAT software accomplishes three basic things:

- NAT can allow you to connect many more machines to the Internet than you have IP addresses for. I first used NAT to connect my home LAN to a dial-up ISP via a single-IP-address PPP connection. (I used the reserved address block 192.168.1.0/24 on my LAN).

- NAT is a good security measure when you use reserved addresses behind the NAT router, since the addresses are not globally routable. It is harder to attack hosts when you can't reach them directly.
- NAT is a good security measure because *no inbound connections are allowed* through the NAT translator unless it is specifically configured to allow them; as we will see, this is a side effect of using NAT software.

I like to refer to NAT routers as "transparent TCP proxy routers." Transparent, because unlike HTTP proxies, NAT routers do not need any configuration nor application software support to work with most TCP-based protocols. NAT routers will proxy outbound connections "automagically."

For every outbound TCP connection, the NAT router intercepts and creates its own TCP connection to the destination host. The NAT router builds a growing list of port translations. Consider two computers that open three TCP connections each to a web server to download the same Web page. At the same time, a Linux workstation opens a Telnet session to rs.internic.net:



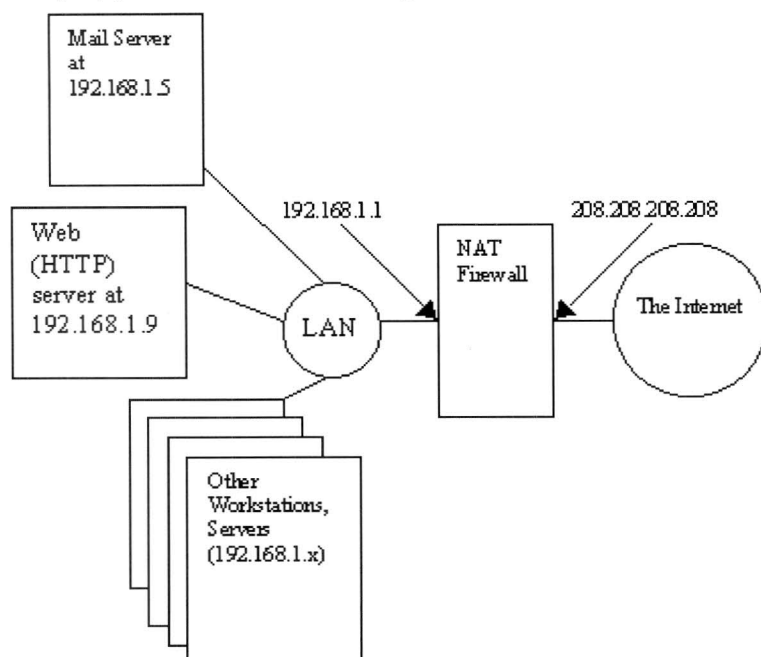
The web server thinks the NAT router at 208.208.208.208 has two browsers running that both just opened the same document and images; the Telnet server thinks that the same computer at 208.208.208.208 opened a Telnet session to it; only the NAT software knows that three computers have seven connections open from behind it.

Transparent NAT works well for TCP connections, but due to the connectionless nature of UDP, NAT works less well for unusual UDP connections (sorry, Quake fans..!)

Since NAT routers *are* hiding many machines behind a single IP address, putting server(s) behind a NAT router becomes a problem, since the NAT software has no way of determining for itself what IP address to forward the inbound connection requests to. This dropping of inbound connections, while allowing outbound connections, makes NAT routers into cost-effective low-end firewalls. Though NAT routers do nothing to prevent users from downloading viruses or trojan horse programs (like the well-publicized trojan horse Back Orifice), but does go a long way toward blocking attempts to connect inbound to the running trojan horse, if accidentally or maliciously installed.

If your NAT router only supports one "real" IP address, you can only have one service on your network listening on the "well known port" for that service; you could have two Web servers

listening on different ports, but not two web servers both listening on (e.g.) 208.208.208.208:80. For example, you have a LAN configured as follows:



This image created using SmartDraw. [Click Here](#) for a free trial copy.

You would configure the NAT software to listen to ports 25 and 80 on 208.208.208.208, and forward connections as follows:

"Listening" Port	"Internal" Address
208.208.208.208:25	192.168.1.5:25
208.208.208.208:80	192.168.1.9:80

If you want to play with NAT software, and you have an old '386 or '486 machine lying around (NAT is easy for routers to do and does not require much in the way of hardware), I recommend IPRoute (which works with any type of network), available at <http://www.mischler.com/iproute/>. Please, read the manual and experiment a bit with the software before sending me questions specific to IPRoute. You'll learn more that way, and I didn't write the program, so I probably shouldn't be the person you talk to for tech support about it, anyway. :-)

Platform Specific Infomation: Note that TCP/IP proxies are not platform-specific. In other words, it works fine to place a MS-DOS based proxy server (such as IPRoute) on a Mac network, or a Linux proxy on a Novell-based IP network. But if you only want to add software, not hardware, to your network, then here are some options I've found. (Note: I do not explicitly endorse the use of any of these products, they're merely listed here for your convenience.)

- **Netware networks:** I understand Border Manager does address translation. Novell's IPX-to-IP gateway (a different product) works for IPX-only networks by tunneling IP sockets through the IPX network to the Netware server, which makes the "real" TCP/IP connection to the destination server. Workstations are protected by virtue of the fact that they're not actually running TCP/IP locally, and don't have IP addresses of their own (they all use the server's IP

address).

- **NT networks:** Microsoft Proxy Server does HTTP proxying, but another solution to consider is WinProxy (<http://www.winproxy.com>) for NAT and HTTP proxy services. I use and recommend WinProxy -- I've never seen NAT software so easy to install.
- **Macintosh networks:** I rarely work with Mac networks. If you know of Mac proxy server, please drop me a line and I'll link to it here.
- **Linux:** Vaguely recent (v1.3.x or greater) Linux kernels include support for "IP Masquerading," which is its name for network address translation. There is a newer kernel option called "transparent proxy" which is not NAT, but rather *forces* all outbound connections to use a proxy server, without the user's knowledge or explicit configuration. Linux goes a step further with Masquerade Loadable Modules (the link, <http://ipmasq.home.ml.org>, is down at the time of this writing; it will probably reappear at a new address), which can explicitly support "wierd" connections such as the seperate UDP connections Quake servers use. See the Linux IP Masquerade mini HOWTO, part of the Linux Documentation Project, at <http://metalab.unc.edu/LDP/HOWTO/mini/IP-Masquerade.html>

12. The Domain Name System (DNS)

The Domain Name System, or DNS, is a service that translates computer names into IP addresses. A name-to-address system is necessary because we humans do not easily remember numbers like, "207.68.156.61", but we can easily remember names like, "www.microsoft.com". The DNS is a hierarchical system, with the top of the system called the "root", and represented by a single period ".". There are twelve (very, very busy) "root" servers on the Internet at the time of this writing. Root servers know where the servers are for the "top-level domains" like .com, .net, .edu, .org, .uk, .de, .nz, .us, and so on.

Let's start with an example: If you ask your local name server for the address of "www.north-america.example.com" the name server will do the following:

1. Check to see if it already knows the address of "www.north-america.example.com" (let's assume it doesn't. The example is more interesting that way.)
2. The DNS server queries a "root" server for the address of "www.north-america.example.com". All fully-functional DNS servers are configured with a static list of root servers, available at <ftp://ftp.rs.internic.net/domain/named.root>.
3. The root server will refer your DNS to a list of ".com" servers.
4. Your DNS will query one of the ".com" servers for the address of "www.north-america.example.com"
5. The ".com" name server queried refers your name server to a list of name servers for "example.com".
6. Your DNS server then asks one of the "example.com" name servers for the address of "www.north-america.example.com".
7. One of two things can happen here. If the "example.com" name server queried knows the address of "www.north-america.example.com" then it returns that address to your DNS server. If the "north-america" subdomain has been delegated to some other name server(s), then that name server list (of name servers that service the zone, "north-america.example.com") will be returned to your DNS, and your DNS will query one of those servers for the address of "www.north-america.example.com".

Note that your DNS remembers, or caches, all the information it retrieves this way. Therefore, if you

asked your local DNS for the address of "ftp.north-america.example.com", then it would directly ask the name server finally referenced in step 7 (above) for the address of "ftp.north-america.example.com". This prevents the top-level and root servers from being more heavily loaded than they already are. (It's also interesting to note that the root servers are also the top-level domain servers for the US domains.) It is possible to set up a caching-only DNS server that processes and caches requests, but isn't directly knowledgeable ("authoritative") about any domains itself.

Domains, Zones, and Authority

There are several different types of name servers. There is one Primary name server for each domain or delegated subdomain ("zone"). A "zone" refers to the domain and subdomain(s) (if any) a server is authoritative for. In many cases "zone" and "domain" mean the same thing, but when you start delegating authority for subdomains, they get their own *zone* to administer, although it's part of your *domain*. For example, the root servers are authoritative for the ".com" *zone* but they aren't authoritative for the entire ".com" *domain*. "example.com" is, in fact, a *subdomain* of the ".com" *domain*, but is a different DNS *zone*. Zone boundaries typically follow administrative control boundaries: since the people managing the ".com" domain are not the same as the people managing the "example.com" domain, a new zone is created and authority for the zone is delegated to that zone's name servers.

Every Primary name server should have at least one Secondary name server. A Secondary name server simply copies the zone information from the zone's Primary server. Secondary name servers also answer DNS requests authoritatively. It is *strongly* suggested that at least one secondary name server be on another physical network. If someone wants to send you mail, and your mail server is unreachable, the mail is queued and retried, but eventually delivered. If the sending mail server is told there is no mail server or host information about your network (which is what happens if all authoritative DNSes are unreachable) then the mail bounces.

If you set up a Primary name server, it is necessary to have the parent domain delegate authority for your zone to you. For example, if you wanted to be the authoritative name server for the domain "reallyslow.net", you would have to ask the administrators for ".net" (InterNIC, in this case) to delegate the zone authority for "reallyslow.net" to you. Similarly, if the engineering department wanted to run their own name server for "engineering.reallyslow.net", then they would have to ask you to delegate the zone "engineering.reallyslow.net" to their name server(s).

It is usually possible to look up an address and come up with a machine name. This is called a "reverse lookup," because instead of getting an address from a name, you are getting a name from an address. The reverse lookup system behaves very similarly to "normal" DNS; in fact, you could almost consider it to be a parallel DNS system. Lookup is done in reverse order by octet with the domain "in-addr.arpa" appended. Let's say you "own" a network 192.168.45.0 with a subnet mask of 255.255.255.0. You would contact the administrator for "168.192.in-addr.arpa" and ask him/her to delegate the authority for the zone "45.168.192.in-addr.arpa" to your name server. On your name server you would create a zone file for reverse lookups that would be authoritative for that zone.

Types of DNS Records

SOA: A Start of Authority record is used at the top of every zone file to indicate the zone that the file is authoritative for. The SOA record also contains administrative contact information, the serial number for the file (which must be incremented whenever the file is updated), and various default timeout and retry values for the domain.

```
reallyslow.net.          IN SOA  turtle.reallyslow.net root.reallyslow.net ([various
```

A: Address records actually provide name-to-address mapping:

```
turtle.reallyslow.net.    IN A      192.168.45.10
```

```
caterpillar.reallyslow.net. IN A      192.168.45.12
```

CNAME: Canonical name records are "alias" records that are often used to map conventional names like "www.reallyslow.net" to the actual name ("A" record) of the computer providing World Wide Web services for the domain. Other names use by convention include "ftp." for ftp services, "mail." for e-mail servers, and "ns" for name servers.

```
www.reallyslow.net.      IN CNAME  turtle.reallyslow.net.
```

```
snail.reallyslow.net.    IN CNAME  caterpillar.reallyslow.net.
```

NS: Name Server records indicate which machines are used as name servers. NS records sometimes point to host names ("A" records), sometimes point to aliases ("CNAME" records), and sometimes just list an IP address.

```
reallyslow.net.          IN NS     turtle.reallyslow.net.
```

```
reallyslow.net.          IN NS     snail.reallyslow.net.
```

MX: Mail eXchanger records indicate which machines are mail servers for a domain and what their preference is. The lower the number, the higher the preference (hey, I didn't invent it.) Other mail servers will try to send mail to the highest preference mail server first. We want email for anyone@reallyslow.net to be delivered to the machine mail.reallyslow.net:

```
reallyslow.net           IN MX 10  mail.reallyslow.net.
```

or, if you used another company to handle your email services...

```
reallyslow.net           IN MX 10  mail.notquitesoslow.net.
```

MX records should not point to CNAME records.

PTR: Reverse lookup pointers are used by the reverse lookup system to map addresses to names (notice the reversed order of the octets:)

```
10.45.168.192.in-addr.arpa. IN PTR    turtle.reallyslow.net.
```

```
12.45.168.192.in-addr.arpa. IN PTR    caterpillar.reallyslow.net.
```

Note that host names never include "@" symbols. An "@" symbol always indicates an email address. The name to the right of the "@" sign is queried for an MX record and mail is delivered to the machine indicated by the MX record(s). In a DNS file, the "@" symbol is a placeholder used to represent "the current domain" as it was named in named.boot. named.boot is the standard file name used by DNS ("named", pronounced "name dee") servers. A basic named.boot looks like this:

```
primary reallyslow.net db.reallyslow.net

primary 0.0.127.IN-ADDR.ARPA db.127.0.0

primary 45.168.192.in-addr.arpa db.inaddr
```

We're telling BIND that it is authoritative for the "standard" zone "reallyslow.net", and also primary for the reverse lookup zones for the subnets 192.168.45.x and 127.0.0.x. (The only entry for 127.0.0.x is 127.0.0.1, which maps to LOCALHOST, which is a reserved address and name for "this machine". In other words, you will always have a VERY fast ping to localhost :-)) The zone file for 45.168.192.in-addr.arpa contains standard PTR records after the SOA record. Note that it's really easy to forget to update named.boot if you add a new domain to your name server (hint, hint.)

If you are going to set up your own name server, I *highly* recommend the book DNS and BIND by Paul Albitz and Cricket Liu (O'Reilly & Associates, ISBN 1-56592-010-4). On the 'net, check out the "BIND Operations Guide" in Windows Write format at <ftp://ftp.software.com/BIND-NT/BOG.wri>.

13. Tips for Building an IP LAN

The part you were waiting for, right?

- First, if you skipped ahead to this section, go back and read the previous sections. No amount of tips will supersede knowing what the heck you're doing.
- If you're not connected to the Internet, and don't already have one or more IP subnets assigned to you, use the addresses reserved for this purpose. They are, 10.x.x.x, 172.16.x.x-172.31.x.x, and 192.168.x.x (see RFC 1597)
- Create a subnet address policy (e.g., .1-.5 reserved for routers, .1 always the default route, .6-.30 reserved for static IP's such as servers, .50-.254 dynamic through bootp/dhcp.)
- Use DHCP or BOOTP to assign workstation addresses. When it comes time to change (after your network has grown a bit), you'll thank me.
- Meticulously track static IP assignments. Create a central database or document listing all static IP's and their associated devices.
- Label router interfaces with their addresses.
- Keep a **current** diagram of your subnets and router connections (include detail on router interface addresses.) If you get into trouble, it'll save you two hours of onsite time if you have to call someone in to help. Personally, I use SmartDraw for this purpose (<http://www.smartdraw.com/>.) Although Visio also works well, SmartDraw (IMHO) is easier to use and the price is right.
- If you have IP-enabled servers, use a firewall. If you are using Windows-based file sharing and have no firewall, use a non-IP protocol to do it (IPX or NetBEUI). You will then need to set either IPX or NetBEUI as your default protocol. Or, get a firewall. IPRoute from Dave

Mischler (<http://www.mischler.com>, \$50US) can be used as an effective, low cost firewall, and it'll run on any '386 or better PC with two NICs; or, if you want to experiment with putting your LAN on the Internet, a low cost and very secure way to do this is with IPRoute and an old '386 PC with a good serial port chip and modem.

- Check out my drawing of an example LAN [here](#).
- More to come / [Your Tip Here](#).

14. WAN Connectivity

This section is in the middle of its first draft. As it's being written, comments and suggestions will help me greatly in making this section as useful as possible for you, the reader.

Concepts

Wide area networking is actually fairly simple conceptually, but can also be one of the most difficult aspects of networking. On the Near Side of the 'Net, however, things rarely get exceedingly difficult, and with a proper understanding, can be quite easy and simple.

It's quite likely you're reading this over a Wide Area Network (WAN) connection-- your dial-up connection to the Internet! What you've done is run a serial cable all the way across town, over streets, under bridges, to your Internet Service Provider (ISP), right?

No?

Oh, I see. You have a serial cable that connects to a modem, that connects (through the phone system) to your ISP's modem, that connects (serially) to a device that allows PPP Internet connections.

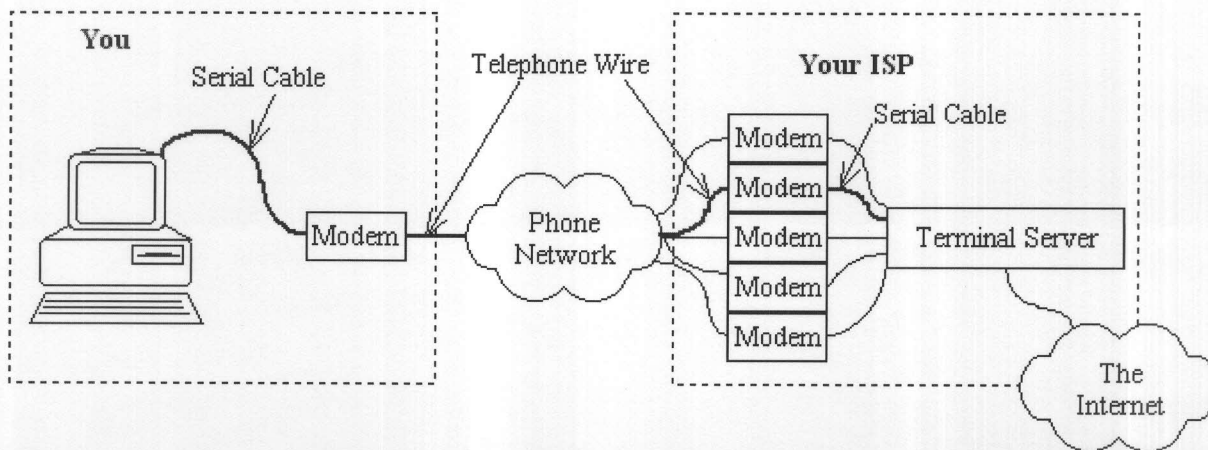


Image created using SmartDraw. [Click Here](#) for a free trial copy.

Follow the transition: Serial Cable, [Modem], Phone System, [Modem], Serial Cable. **From this perspective, a modem is nothing more than a serial cable extender, that allows you to run a serial cable through the phone network.** And that's all it is-- a serial cable extender. From a layer

1/layer 2 perspective, the sole function of a modem is to allow you to extend your serial connection through a phone system. **Most WAN links are simply some method of serially connecting two routers through the public telephone network.** The only real differences are in speed and flexibility.

Ok, let's cover some terminology:

Point-To-Multipoint Networking

Networking where one device may be physically connected to multiple devices, such as when using Ethernet or Token-Ring. A layer two address (typically, a MAC address) is required to indicate to the network which device you're talking to. Typically used for LAN connectivity.

Point-To-Point Networking

Networking where one device is physically connected to one device, such as when using a serial cable (or extended serial cable) and the PPP protocol. There is no concept of MAC address in this case (which can present some difficulties when routing IPX over WAN links, but that's outside the scope of this document.) Typically used for WAN connections.

PPP

The Point to Point Protocol. Provides a standard way of running multiple protocols simultaneously over a WAN link.

SLIP

The Serial Line Internet Protocol. Provides a way of running IP over a dial-up WAN link. Only occasionally still found in use, it has been largely replaced by the more flexible PPP.

Modem

Provides a means of extending a [digital] serial link over an [analog] voice network.

ISDN

Integrated Services Digital Network. Originally designed to replace the Plain Old Telephone System (POTS), high price and restricted availability have restricted its adoption primarily to medium-speed WAN connections. More on ISDN in a bit.

Frame Relay

A point-to-point, point-to-multipoint hybrid that allows multiple "virtual" connections, or circuits, to exist on a single physical connection. A frame-relay "cloud" in the center, managed by the intermediary telco(s), manages the frame-relay network so you don't have to. Or, that's the way it's supposed to work, anyway. :-)

Frame Relay PVC

A **P**ermanent **V**irtual point-to-point **C**ircuit through the frame-relay cloud.

POP

Point of Presence. Typically used to describe the a location from which a service is provided. For example, a ISP modem bank can be referred to as a modem POP, or a frame-relay switch can be referred to as a frame-relay POP.

56k

A digital WAN circuit leased from the phone company. Allows communication at 56kbps bidirectionally. Can be connected directly to another office location, or to the nearest frame-relay POP. Requires a 56k CSU/DSU to be useful as a digital WAN link.

T1

A digital WAN circuit leased from the phone company. Originally designed to reduce the need for copper under streets (they were running out of room,) a T1 is configured into 24 digital channels, each of which can carry one digitally encoded voice conversation. For use as a serial cable extender (WAN link), a T1 CSU/DSU is required.

CSU/DSU

Converts the digital signaling of a serial cable to the digital signaling of the telco network;

functionally, the same role as a modem. T1 CSU/DSU's also handle the T1 channelization, which is why they're much more expensive than 56k CSU/DSU's. Conceptually, a CSU/DSU is two devices rolled up into one: a Channel Service Unit, which handles telco signaling; and a Data Service Unit, which converts the serial cable signaling into one or more sets of signals the CSU can easily deal with.

ISDN

ISDN, or Integrated Services Digital Network, was the digital technology that was supposed to replace analog telephones. However, lackluster (and 'lackluster' is being generous) support from US phone companies have hobbled ISDN's chances of ever replacing the current analog networks. Phone conversations are typically analog between you and the local phone switch; digital from switch to switch; then analog from the destination switch to the other person you're talking to. This analog-digital-analog conversion makes the engineers of modem manufacturers lose sleep. Since ISDN is end-to-end digital, it is well suited to carrying data as well as voice. The basic consumer ISDN connection is a Basic Rate Interface, or BRI, circuit. A BRI is physically installed as a single pair of copper wire, but has three logical "channels" (think TV channels.) These channels are referred to as "B" or "D" channels. BRI "D" channels are 16kbps and are used by ISDN equipment for talking to the telco switch ("You have an incoming call" or "I want to call this number"). ISDN "B" channels are 64kbps and a BRI circuit contains two of them. For this reason, people often refer to BRI as "2B+D". Each "B" channel is considered to be a separate phone line by the phone company, which becomes important if you want to use both of them simultaneously for dial-up connectivity, or when the per-minute bill arrives from the phone company.

That connection from a single pair of copper is known in ISDN circles as a "U" interface, and the phone company expects you to attach an "NT1" to it. An NT1 then provides two-pair "ST" interfaces to the various ISDN devices around your house. In practical use, most people don't use ISDN for voice. Hardware manufacturers have picked up on that fact and will usually build the NT1 right into the device-- the device, then, is said to have a "built-in NT1" or have a "U Interface". Devices that expect an external NT1 are usually described as having an "ST" interface and are less expensive than their NT1 Interface counterparts. In most cases, when using ISDN for networking purposes, you will want to purchase a device with a built-in NT1.

Submitted by Tony F.:

In Europe connection to the ISDN network is via the 'S' interface. The difference being (in no technical terms) is that the conversion from the signals on the copper...to digital format is done by the service provider. In the US the ISDN device that you buy does this bit as well. ie the customer pays [for and owns] the conversion [hardware].

Multilink PPP and BACP

Although ISDN is split into two channels, dialing two separate [regular] PPP connections to an ISP is not desirable; you would have two different IP addresses, and the best throughput possible in either direction is 64kbps (sending data on one channel while receiving data on the other.) Since most "near side of the 'net" connections are primarily receiving data, having the ability to mostly receive data is important. Enter Multilink PPP (MLPPP). Simply put, Multilink PPP allows a single logical PPP connection to span multiple physical connections. A newer protocol, the Bandwidth Allocation Control Protocol (BACP), allows channels to be added and dropped dynamically, typically in

response to higher utilization. Typically, MLPPP asks for two phone numbers to dial, but the two phone numbers are usually identical. BACP will usually ask for the minimum and maximum number of channels to connect. The minimum is generally 0 for outbound-only Internet connections, and 1 for "listening" connections such as mail or Web servers; the maximum number of channels is almost always 2.

See Dan Kegel's ISDN Page for much more ISDN information.

<http://www.alumni.caltech.edu/~dank/isdn/>

56k Connections, Analog

Ahh, the wonderful, ubiquitous 56k dialup connection. It's often all that's required for a small LAN to send and receive e-mail and do some light Web browsing. Many ISPs only offer one POP e-mail account for a dialup connection, but there are other services you can use to add POP accounts for free: Hotmail and Yahoo come to mind. Because you generally get one (varying) IP address, some means of "hiding", or proxying, several "fake" IP addresses behind your single "real" IP address. See the section on NAT for more information. Personally, I use WinRoute for this purpose. One of these days, when I have the spare time (yeah-- right), I intend to get my Linux machine doing this via IP Masquerading. However, my current configuration is working quite well, and fixing what's not broken tends to rank fairly low on the priority list.

56k Connections, Digital

Coming soon..?

T1 Connections

Coming soon..?

Frame-Relay

Coming soon..?

Routing Over WAN Links

Coming soon..?

[Hmm.... this looks like it's going to be a long section.]

15. Update Notifications/Comment Form

Notification Subscription and/or Comments

I do update Daryl's TCP/IP Primer on an irregular basis; if you'd like to be notified of these updates, or if you want to send me a comment or suggestion, use the appropriate boxes below. I will not use your name nor email address for any other purpose than to alert you of updates to Daryl's TCP/IP Primer, and those notifications (even for minor updates) don't happen very often; expect an email every 2-6 weeks or so (irregularly) for minor updates; major updates anywhere from quarterly to

annually. I add Q&A's as time allots; possibly two in one day and none for the month following. You can always change your notification options later by resubmitting the form (instructions will be provided in every email.)

I've been considering breaking the document into smaller sections that can be more easily downloaded (but less easily printed.) If you have an opinion on this, please include as a comment.

Finally, if you do submit the comment form more than once, you will not receive duplicate notifications; you can, in fact, change your notification option this way.

Notification and Comment Form	
*Your Name:	<input type="text"/>
*Your E-Mail Address:	<input type="text"/>
Notify Options:	<input checked="" type="checkbox"/> Major Updates (e.g., new chapter; rare) <input checked="" type="checkbox"/> Minor Updates (content added; less rare) <input checked="" type="checkbox"/> New Q&A Added (every week or so, so far)
Comments:	<div><div></div></div>
<input type="button" value="Submit Query"/>	

16. Questions and Answers

The following are questions submitted to me via e-mail. The answers may not always be complete, and quite often there are unmentioned exceptions (that, of course, prove the rule :-)

As usual, use any information here at your own risk; I am not responsible if any errors or omissions that adversely affect you.

If you submit a question to me, please include whatever details you can to help me answer. I don't guarantee a response; if I do respond, I may post the response here, without your full name, edited for brevity, and after altering any IP addresses to preserve your anonymity.

Question added 6/5/1999, submitted by Kent

Q: This is a great site. Thank You. [You're welcome.] I do have one question concerning subnetting and when to do so. How many nodes can you put on one TCP/IP subnet before it requires segmenting your network? I am referring to a Lan with approx. 300 users. Is there a reason why I can't use a standard 255.255.000.000 subnet. I will only be assigning addresses in my DHCP scope as the network requires them.

A: This is a good question, and really is more of a layer two question than a TCP/IP question. I would not run a 300 user lan on a single 10Mbps Ethernet segment; however, I wouldn't balk at a 300 user network segmented into 12 or 24 switched partitions using a centralized Ethernet switch. So the real question here is, "will my current layer two network topology support 300 users on one segment?" You can put as many nodes as you want on one TCP/IP segment; however, that lack of limitation does not apply to Ethernet. (I would ensure no Windows boxes are running NetBEUI, though.)

Remember, a switch "segments" networks on layer two, and a router "segments" on layer three. The main difference, from a topology planning standpoint, is that switches forward broadcast packets and routers don't. Thus, switching becomes a problem quickly with "loud" protocols like NetBEUI, since switching doesn't reduce or segment broadcast traffic.

You can use a subnet mask of 255.255.0.0 to put up to 65,534 hosts on a single routed network segment; or you can use a subnet mask of 255.255.254.0 to put up to 512 hosts on a network segment. I'm assuming you're using "reserved" addresses (such as 10.1.x.x) behind a NAT firewall or proxy, so the choice of subnet mask is yours. The choice of whether or not to segment by switching or routing is also yours; I tend to prefer switching, since it tends to keep things simpler.

Question added 1/23/1999, submitted by NBK

Q: How vulnerable is Linux against Net attacks compared to NT??? Damn NT has to many holes....

A: In both cases, it depends on the administrator :-)

a good packet filter or (better yet) firewall, good knowledge of the security issues of the services the box is providing, and keeping current on the security updates/ mailing lists for the OS'es and running services makes for a pretty strong box. Any badly installed service can present the opportunity for a full breach; be sure to read the security FAQ's (and I'll often scan cracker websites) for the OS and the services you're making available to the public.

Question added 12/3/1998, submitted by David

**Q: This is to request from you a tutorial on TCP/IP.
Thank you very much.**

[Answer: can you be more specific? Platform, etc?]

Actually I'm looking for an overview on the internet network. How the providers build their network...

How do they get inteconnections...

What are the critical economical issues for internet on the next years...etc

A: Hm... That's intentionally outside the scope of the Primer (hence the subtitle, "...the near side of the 'net.") For the information you're looking for, search for "BGP4" re: interconnections, and regarding economic issues (etc) try any of the Internet trade rags for the professional pundits :-)

Doing generic dialup and hosting does not (IMHO) have an entry level any more; the services are very commoditized and the economies of scale involved will squeeze out the smaller non-value-added providers. But (apologies to Dennis Miller) that's just my opinion, I could be wrong.

Question added 10/12/1998, submitted by Joanne

Q: The part I don't understand is: what is the reason to subnet? You can't possibly get more destinations that way, I mean, 32 bits are 32 bits. There's only 4 billion possible internet destinations, no matter how you split it up. So what does subnetting do for ya?

A: Subnetting does two things, depending on what context you're in:

If you're a workstation (or server), the subnet mask is used to determine whether the destination IP address is on your same subnet; if so, the workstation will attempt to ARP the destination's Ethernet card address and deliver the data directly; remember, **the first routing decision is made by the workstation**, and the decision is: whether or not to send the packet to a router.

Routers keep their routing tables manageable by clumping large blocks of addresses together using broad subnet masks ("Network Prefixes"). In the old days of classful routing, routers would have to keep track of each "Class C" address individually, which was causing extreme growth of routing tables; CIDR routing allows you to clump as many "Class C" networks together as you want (in powers of 2.)

So, you may ask, what about servers that also act as routers? In which category to they fall? Well, I lied when I said that subnetting does different things depending on context; it's just that most IP end stations (workstations) don't bother trying to keep track of the whole network; they just know that "these addresses are local, and I'll send anything else to my default gateway/router."

Question added 10/6/1998, submitted by Bob

Q: Is it possible with IE or netscape to address a web server by its MAC address?

A: It sounds like you're asking if you could run HTTP over DLC; the short answer is "no."

The long answer: the HTTP protocol is based on the TCP protocol, which is based on IP; therefore, both the client and server must already be running IP for HTTP to work. You could force client and server IP address into their local ARP caches if they are on the same subnetwork (bounded by routers), but I don't know how well that would work (I doubt the IP stack checks its arp cache before it determines whether or not a given IP is on a locally attached subnet.) If it did work, you could then type the (fake?) IP address of the server into your browser's location line to pull pages. The server would then reply to your (fake?) IP address.

Alternatively, if there is an IP router involved, you could play with its ARP cache; routers are more likely to be forgiving about having multiple IP subnets (or, network prefixes, in RFC 1812 parlance) on the same subnetwork than, say, Win95 workstations.

Note that on any point-to-multipoint network (like Ethernet or Token Ring, but not including serial PPP or HDLC connections), the most basic address (in the layer 2 MAC header) is the MAC address. But you cannot type a MAC address into 'IE or netscape' and connect to a web server; even if you could, the web server would not know what IP address nor TCP socket number to reply to.

Question added 9/30/1998, submitted by Jim

Q: I just have a quick question, its regarding Windows 95 (Yeah, I hear you screaming), when you set the computer to 'disable DNS', and don't set a gateway address (all via control panel) and disable WINS--how is anything assigned to the computer? Is it fair to assume its BOOTP, or something else?

A: Probably DHCP;

BOOTP assigns the IP address, subnet mask, default gateway (route), and (if memory serves) the DNS information. DHCP allows for a bunch of other information to be sent to the workstation, including WINS server addresses. DHCP also has a facility for "lease expiration", where addresses that are not renewed are returned to the pool of available addresses; under BOOTP, IP addresses are permanently associated with the NIC's MAC address, so if you throw out the NIC, the IP address is "lost." Win95 does not support BOOTP.

DHCP and WINS are two very different things, they just seemed to "appear" at the same time (with the introduction and subsequent popularity of Windows 95 and NT Server 3.5x). DHCP is used for automatically configuring workstations with all the information they need to access the TCP/IP resources available to them, including IP address, subnet mask, default gateway, and on Windows NT networks, WINS server addresses. WINS is like DNS for NT networks; WINS is used to "advertise" and locate NT server and (win95|nt) workstation resources on the NT network, such as shared drives and printers. DHCP is a non-Microsoft-specific "upgrade" to BOOTP, WINS can be described as a Microsoft Networking version of DNS. (Novell's version of WINS for distributing SAP information is called DSS, or Domain SAP Server.)

BTW-- Win95 doesn't make me scream, but don't bring any Win3.X machines by unless you're equipped with earplugs :-)

17. Other Sources

On the 'net: (in no particular order)

Uri's TCP/IP Resources List: 'This posting contains a list of various resources (books, web sites, FAQs, newsgroups, and useful net techniques) intended to help a newbie to learn about the TCP/IP suite of protocols.'

The NT Shop: Information and links regarding Internet Security, specifically Internet Security as relates to Windows NT systems.

Patrick's MCSE Place: Lots of MCSE stuff, links.

Alliance Datacom Frame Relay Tutorials: Much information about Frame Relay and related technologies. (plus Links Galore!)

Your link here.

In Print: (with convenient links to purchase the books from Amazon.com)

This is a look at my bookshelf- I have included all of the books with more than one crease in the binding.



Topic	Recommendation
Setting up UNIX services for the Internet, including details on TCP/IP and TCP/IP services like DNS and SENDMAIL. Helpful, even if you never touch UNIX.	<u>TCP/IP Network Administration</u> by Craig Hunt Published by O'Reilly and Associates, Inc.
Everything you ever needed to know about DNS. Referred to as the "DNS Bible" or simply "The DNS Book" on DNS mailing lists.	<u>DNS and BIND</u> by Paul Albitz, Cricket Liu, Mike Loukides Published by O'Reilly and Associates, Inc.
Packet filtering and general Internet security	<u>Building Internet Firewalls</u> by D. Brent Chapman, Elizabeth D. Zwicky, Deborah Russell Published by O'Reilly and Associates, Inc.
UNIX System Administration. Covers issues with several *nix flavors.	<u>Essential System Administration : Help for Unix System Administrators</u> by AEleen Frisch Published by O'Reilly and Associates, Inc.
My Linux command reference. Always at my elbow when I'm doing anything interesting on my Linux box.	<u>Linux in a Nutshell</u> by Jessica Perry Hekman, Andy Oram (Ed) Published by O'Reilly and Associates, Inc.
High speed Internet connectivity	<u>Getting Connected : The Internet at 56K and Up</u> by Kevin Dowd, Mike Loukides Published by O'Reilly and Associates, Inc.
JavaScript: This book had everything I needed to make the Subnet Calculator work.	<u>Javascript : The Definitive Guide</u> by David Flanagan Published by O'Reilly and Associates, Inc.
Interdomain (Backbone) routing-- how the big boys do it.	<u>Internet Routing Architectures</u> by Bassam Halabi Published by Cisco Press
MS SQL Server administration. My most worn SQL Server book.	<u>Using Microsoft Sql Server 6.5</u> by Stephen Wynnkoop Published by Que Education and Training

In fact, every book I've read from O'Reilly and Associates has been very good. If you see one of their books relating to a subject you're interested in, my advice is to buy it.

18. Glossary

"A" Record

A DNS host record, used to name-resolve all non-email addresses.

Address Lookup

see Domain Name Service.

ARP

Address Resolution Protocol. On LANs, this is used to get the Layer 2 address of a host, so that IP transmission can take place over Layer 2 protocols like Ethernet or Token-Ring.

Banana

~~A tropical fruit.~~

Contributed by Brent G. Gratias:

banana (be-nàn'e), name for a family of tropical herbs (the Musaceae), for a genus (Musa) of herbaceous plants, and for the fruits they produce. Bananas are probably native to tropical Asia, but are widely cultivated. They are related to the economically valuable MANILA HEMP and to the BIRD-OF-PARADISE FLOWER. Banana plants have a palmlike aspect and large leaves, the overlapping bases of which form the so-called false trunk. Only female flowers develop into the banana fruit (botanically, a berry), each plant bearing fruit only once. The seeds are sterile; propagation is through shoots from the rhizomes. Bananas are an important food staple in the tropics. The Concise Columbia Encyclopedia is licensed from Columbia University Press. Copyright © 1991 by Columbia University Press. All rights reserved. [some people take humor so *seriously*!]

Class

see "Class A", etc. in "IP Addresses, Subnet Masks, and Subnetting, Part A," above.

DNS

see Domain Name Service.

Domain Name Service

All communication on the Internet is done based on IP addresses. Name Service allows "us humans" to use names for services, which, for us, are much easier to understand. The "address lookup" process converts address names like "www.microsoft.com" to IP addresses like "207.68.137.9".

Host

For purposes of this discussion, a Host is an IP-aware machine connected to an IP network. Although a host can have more than one interface, those hosts usually perform routing functions, and are therefore called routers when referred to specifically.

Interface

A connection to a network. Usually either a network card or a serial WAN link.

Internet Protocol

The means of communication on the Internet, usually abbreviated to IP. IP involves four-byte address, where each byte is expressed in decimal numbers and separated by a period, like "168.192.1.1".

Internet Service Provider

The people you pay to get Internet access. (Not the phone company.)

Internetwork Packet eXchange (IPX)

A protocol stack typically used on Novell networks. Useful for its simplicity of configuration, this protocol (in its current implementation) does not scale real well to large networks, for reasons touched on in section 3.

IP

see Internet Protocol.

IPX

see Internetwork Packet eXchange (IPX).

ISP

see Internet Service Provider.

Layer 1

In OSI Model terms, the conceptual networking layer that defines electrical signaling on a wire.

Layer 2

In OSI Model terms, the conceptual networking layer that defines physical addressing and packetizing over a given wire.

Layer 3

In OSI Model terms, the conceptual networking layer that defines logical addressing and routing.

"MX" Record

A DNS Mail eXchanger record, used to name-resolve email servers for purposes of mail delivery.

Name Resolution

see Domain Name Service.

Network

see Segment

Network Mask

Under RFC 950, this was the fixed part of the "Subnet Mask" that was determined by the class of address.

Network Prefix

Under RFC 1812, "Network Prefix" combines and replaces "Network Mask" and "Subnet Mask," both RFC 950 concepts.

Protocol Stack

The software used to communicate on a network using a given protocol, e.g., "I'm running Windows for Workgroups and using the Microsoft TCPIP-32 implementation of the TCPIP protocol stack." The word "stack" is derived from the layered nature of networking.

Router

A host, with multiple interfaces, that performs routing.

Routing

The process of intelligently forwarding packets from network segment to network segment.

Segment

In this discussion, a "network segment" is a collection of hosts and/or layer 2 networking devices bounded by routers.

TCP/IP

Transmission Control Protocol/Internet Protocol: TCP is a Layer 4 protocol not covered by this document. See Internet Protocol, Protocol Stack.

WAN

see Wide Area Network

Wide Area Network

A term generally used to describe any network that includes at least one dedicated link that involves paying the phone company. (Specifically, WAN should only be used when the link takes you from one city to another, but then who's watching?)
